

Preparing to Run on Aitken Rome Nodes

To help you prepare for running jobs on Aitken's Rome nodes, this short user guide includes information on the general configuration of Rome nodes, compiling your code, and running PBS jobs.

Overview of Aitken Rome Nodes

Aitken includes 1,536 Rome nodes, which are partitioned into twelve physical racks. Each node contains two 64-core AMD EPYC 7742 Rome sockets (2.25 GHz) and 512 GB of memory. The nodes are connected to the Aitken InfiniBand network (ib0 and ib1) via four-lane High Data Rate (4X HDR) devices and four-lane High Data Rate (4X HDR) switches for internode communication. The ib1 fabric is used mainly for I/O and is connected to the Pleiades Lustre filesystems. In addition, Aitken, Electra, and Pleiades share the same home filesystems, Pleiades front-end systems (PFEs), and PBS server. You can access Aitken only through PBS jobs submitted from the PFEs.

Operating System and Software Modules

As of Dec. 20, 2021, the operating system on all Rome compute nodes is the Red Hat Enterprise Linux-based Tri-Lab Operating System Stack Version 3 (TOSS 3). Batch jobs use `:aoe=toss3` by default, so you do not need to add `:aoe=toss3` in your PBS resource request. The default `$MODULEPATH` includes:

- /usr/share/modules/modulefiles
- /nasa/modulefiles/toss3
- /nasa/modulefiles/spack/gcc-4.8/
- /nasa/modulefiles/pkgsrc/toss3/

If you include `:aoe=sles15` in your resource request, PBS will provision the requested number of Rome compute nodes into the SLES 15 SP2 operating system, and the `$MODULEPATH` will include:

- /nasa/modulefiles/sles15
- /nasa/modulefiles/pkgsrc/sles15

Note: PBS provisioning can delay the start of a job; it may require rebooting compute nodes into the requested operating environment `aoe`.

Compilers for Rome Nodes

There are several compilers you can use to build your application:

- Intel Compilers:

Some Intel compiler modules are available in the `/nasa/modulefiles/toss3` directory, however, several are in the `/nasa/modulefiles/testing` directory. To see what versions are available in these two directories, run:

```
module use nasamodulefilestesting
module avail compintel
```

Existing executables that are built with support for **AVX512** (for running on Skylake and Cascade Lake processors) will not run on the Rome nodes. Others that are built with support for **AVX2** may run on the Rome nodes; however, you should always check the result of your runs using existing executables. To rebuild your application using Intel compilers, use **-march=core-avx2**.

Note: Although AMD recommends using **-xCORE-AVX2**, we do *not* recommend using it, as the resulting executable will run but will not actually take advantage of **AVX2** instructions. Executables built with **-xCORE-AVX2**, **-xSSE4.2**, etc., will not run on Rome processors.

- **GNU Compilers:**

The default version of GNU Compiler Collection (GCC) under TOSS 3 (/usr/bin/gcc) is GCC 4.8.5. However, support for the Rome architecture is only available with GCC 9.x and later. There are several GCC versions in the /nasa/pkgsrc/toss3 directory that can be used with the Rome nodes. For example:

```
module avail gcc
-- /nasa/modulefiles/pkgsrc/toss3/ --
gcc/10.2 gcc/10.3 gcc/7.5 gcc/9.3
module load gcc/10.3
```

The recommended compiler flag to use with GCC is **-march=znver2**.

- **AMD Optimizing C/C++ Compilers (AOCC)**

To see available versions before loading an AOCC compiler, use the **module avail comp-aocc** command. For example:

```
module avail comp-aocc
-- /nasa/modulefiles/toss3 --
comp-aocc/3.1.0

module load comp-aocc/3.1.0
```

The recommended compiler flag to use with AOCC is **-march=znver2**.

- **PGI Compilers:**

Some PGI compiler modules are available in the /nasa/modulefiles/testing directory. There is no support specifically for Rome processors. If you want to compile your code using a PGI compiler, the recommended compiler flag is **-tp=zen**

For more compiler options, see [Compiler Options Quick Reference Guide for AMD EPYC 7xx2 Series Processors](#).

MPI Library for Rome Nodes

HPE MPT is the only recommended MPI library for use on the AMD Rome nodes under TOSS 3. Use the **mpi-hpe/mpt.2.25** modulefile in the /nasa/modulefiles/toss3 directory. Earlier versions such as **mpi-hpe/mpt.2.23** do not work properly for the Rome nodes.

```
module load mpiâ hpe/mpt.2.25
```

Note: Although there is more than one version of MPT library installed on NAS systems, **mpiâ hpe/mpt.2.25** is currently the only reliable version for running on the Rome nodes. You can access this version by loading **mpiâ hpe/mpt**, which points to the recommended version and environmental settings:

```
% module load mpi-hpe/mpt
```

Running OpenMP or MPI/OpenMPI Hybrid Codes

Each Rome node has 128 physical cores. When Simultaneous Multi-Threading (SMT) is enabled in BIOS, there are 256 logical cores for use by OpenMP or MPI/OpenMPI hybrid codes.

To check whether SMT is enabled, run the following command on a Rome compute node:

```
proccpuinfo processor
```

If it shows 256, then SMT is enabled.

We recommend using the [mbind.x tool](#) for process and thread binding, as shown in the following example with four MPI processes and 32 OpenMP threads per process:

```
mpiexec n uscicontoolsbinmbindx t aout
```

If you use HPE MPT, you can also use the [omplace tool](#) for pinning MPI/OpenMP hybrid code.

Running PBS Jobs on Aitken Rome Nodes

To request Aitken Rome nodes, use `:model=rom_ait` in your PBS script, as shown in the example below.

The `devel`, `normal`, `long`, `debug`, and `low` queues are available for use with Rome nodes.

Note: The Standard Billing Unit (SBU) rate for Rome processors is 4.06.

Sample PBS Script For Aitken Rome Nodes

```
#PBS -l select=2:ncpus=128:mpiprocs=128:model=rom_ait
#PBS -l walltime=8:00:00
module load mpihpempt
module load compintel
```

```
mpiexec np aout
```

For more information, see [AMD Rome Processors](#).

Article ID: 657

Last updated: 20 Dec, 2021

Revision: 56

Running Jobs with PBS -> PBS on Aitken -> Preparing to Run on Aitken Rome Nodes

<https://www.nas.nasa.gov/hecc/support/kb/entry/657/>